Label Visualization and Exploration in IR

Omar Alonso
Microsoft
omalonso@microsoft.com

arXiv:1612.03316v1 [cs.IR] 10 Dec 2016

ABSTRACT

There is a renaissance in visual analytics systems for data analysis and sharing, in particular, in the current wave of big data applications. We introduce RAVE, a prototype that automates the generation of an interface that uses facets and visualization techniques for exploring and analyzing relevance assessments data sets collected via crowdsourcing. We present a technical description of the main components and demonstrate its use.

1. INTRODUCTION

The adoption of visual analytics in the information retrieval (IR) community has been relatively low compared to other areas in computer science, like databases, that process massive large data sets. Visual analytics, a combination of techniques drawn from information visualization, data mining and statistics, allow users to directly interact with the information presented to gain insights, deduce conclusions, and make better decisions.

With the ever increasing number of data sets and emerging new data sources, the combination of automatic analysis and visual tools offers potential to gain better understanding on the many assets that the typical IR researcher, data analyst or relevance engineer has to deal with on a daily basis. Furthermore, the inclusion of crowd data in the form of labels that can be used for training machine learning models or evaluating the quality of search engine results, offers new opportunities for using visualization techniques.

In IR, collecting high quality document relevance assessments pairs is a crucial step for building relevance models. The task, which is very subjective, consists of assessing the relevance of a document to a given topic or query. The human (editor, judge, or worker) performs a visual inspection of the document and provides a label on a particular relevance scale. Finally, label quality control mechanisms are used to produce the final data set.

In contrast to infographics or static visualization tools that can present different types of metrics and summary statistics, we are interested in interactive visualization and exploration of preferences or labels from the crowd in the context of a particular IR experiment that contains a human intelligence task. That is, visual exploration of workers' preferences for certain scenarios as well as other behavioral data. In other words, visual analytics for data exploration where humans also data providers along with data about their work performance.

We argue that visual analytics techniques should be part of the relevance assessment process to help gather better labels and discover potential issues in the experimental design or worker performance by looking at the entire data set [1]. Visualization tools help users in situations where seeing the structure of a data set in detail is better that seeing only a brief summary [7].

Instead of focusing on the visualization of search results for a query or other descriptive statistics, we are interested in exploring the output of a relevance assessment task with the goal of recognizing patterns. Why do certain workers disagree on specific documents or topics? Are there any relevance cues on the presentation that may confuse a worker? Can we identify difficult tasks?

In industrial settings, thousands of thousands of labels are collected weekly using data pipelines that select query document pairs and upload them to internal or external crowdsourcing platforms. The data analyst then looks for specific metrics or anomalies in data sets using traditional database-driven tools. In this context, visual analytics should allow users to analyze data when they do not know exactly what questions they need to ask in advance.

We automate the construction of a visualization interface given the results of a crowdsourcing task. RAVE (Relevance Assessments and Visual Exploration) is a prototype that uses as input the results of a labeling task and produces a faceted-based visualization interface for exploring and analyzing relevance assessments. We demonstrate how RAVE can be utilized to gather better insights from judges, data sets, and labels.

2. SYSTEM OVERVIEW

We make the following assumptions in our architecture in terms of tools and data access. Our user, the data analyst, has access to a database of queries, topics, and documents. Human intelligence tasks are implemented in an external crowdsourcing platform (e.g., Mechanical Turk, CrowdFlower, etc.) or internal equivalent tool.

The user begins the implementation of the experiment by sampling query-document pairs $\langle q, d \rangle$ from a database. The second step is to annotate the sampled data by running some classifiers and NLP tools such as query type identification (e.g., navigational, informational, transactional) and named-entity recognition (e.g., person, organization, location, etc.) to augment the original data with annotations $\langle e_1, \ldots, e_n \rangle$. Once the crowdsourcing task is completed, the labels provided by workers and other assignment metadata (e.g., worker id, time spent, approval rate, etc.), $\langle l, a_1, \ldots, a_m \rangle$, are available. We can think of the underlying data representation as query-document pairs with query annotations, labels, and assignment metadata. That is, for a single assessment, a tuple

ACM ISBN 978-1-4503-2138-9.
DOI: 10.1145/1235

Query	doc_A	doc_B	query length	query type	has_entity	label	worker_id	work time
youtube	r1	r2	1	navigational	company	A	1	19
youtube	r2	r1	1	navigational	company	A	2	7
youtube	r1	r2	1	navigational	company	A	3	8
selena gomez	r2	r1	2	informational	person	same	1	21
selena gomez	r1	r2	2	informational	person	B	4	37
selena gomez	r2	r1	2	informational	person	B	3	9

Table 1: Data description example for a relevance assessment task. Columns doc_A and doc_B represent the content of the A-B comparison; r1 and r2 are the rankers. Query length in characters, query type, and has_entity are query annotations. Worker_id and work time in seconds represent assignment metadata.

$\langle q, d, e_1, \ldots, e_n, l, a_1, \ldots, a_m \rangle$ where l represents the label. Table 1 shows an example.

Relevance assessment is a visual exercise and capturing the image of what the worker sees at assessing time is an important part of our approach. Each document is saved as an image and the workers perform the task looking at the same set of images. This allows our user to see the same content as workers.

3. RELEVANCE ASSESSMENT VISUALIZATION AND EXPLORATION

As mentioned earlier, RAVE uses the image of the document as the visual focus and query annotations and assignment metadata as facets. The prototype automatically generates visualizations and facets for a couple of available tools: Pivot[1] and Exhibit[2]. We now describe the specific details for the automation.

3.1 Human Intelligence Task

As driving example, we would like to evaluate the quality of a new ranking function against an existing baseline. That is, assess the relevance of two ranking functions r_1 and r_2, each returning a fixed number of documents $d_1, \ldots, d_n$ (where $n < 10$) as results for a query q. For collecting the assessments, we create an A-B comparison task that shows the query and the results for the two rankers in random order (a ranker may appear in column A or B). The task for the workers is to select which search results they prefer according to three choices: A is better, B is better, or they are the same as third option.

As part of the data preparation step, the tool captures a debranded SERP (Search Engine Results Page) screenshot for a query document pair. The document, in this case the ranked hit list, is saved as an image using standard libraries. A debranded page means that there are no specific user interface items that may bias workers.

A bit more processing is needed for Pivot, which requires the use of the Pauthor command line tool for generating Deep Zoom images. For Exhibit, the thumbnail version of the original image is also produced. A configuration file is needed for specifying which columns from the results data file corresponds to which facets.

3.2 Generating a Pivot Collection

Pivot collections are stored in a CXML schema that defines facets and other properties. In essence, the collection is a `cxml` file that describes the facets and contains all the elements needed for presentation. The generation of the collection works as follows. The code first outputs the facets that we are interested in (`FacetCategory`) and then loops through the rows of the input file (the experiment task results) and outputs an item for each entry (`Item`).

[1]http://bit.ly/1DsLdC6

[2]http://simile-widgets.org/exhibit/

```
<FacetCategories>
  <FacetCategory Name="Worker" Type="String"/>
  <FacetCategory Name="Query" Type="String"/>
  ...
  <FacetCategory Name="Answer" Type="String"/>
</FacetCategories>
<Items>
  <Item Id="0" Name="..." Img="Selena_GomezB.png">
   <Description>...</Description>
   <Facets>
    <Facet Name="Worker">
       <String Value="4"/>
    </Facet>
    <Facet Name="Query">
       <String Value="Selena Gomez"/>
    </Facet>
    <Facet Name="Answer">
      <String Value="B"/>
      </Facet>
   </Facets>
  </Item>
  <Item Id="1" Name="..." Img="Selena_GomezB.png">
   <Description>...</Description>
   <Facets>
    <Facet Name="Worker"><String Value="1"/>
    </Facet>
    <Facet Name="Query">
       <String Value="Selena Gomez"/>
    </Facet>
    <Facet Name="Answer"><String Value="Same"/>
     </Facet>
   </Facets>
</Item> ... </Items>
```

Figure 1: Snippet of the cxml code that describes the collection.

The final result is an XML file with all the facets, values and images that can be visualized. To be able to see the visualization, a Silverlight plug-in or the Pivot viewer application are needed for rendering the `cxml` file.

Figure 1 shows a snippet of the schema and collection code. Figure 2 shows the exploration of the results sets for a specific query. We can observe visually that ranker that appears in column A wins over ranker in column B and that a few judges believe that both are the same. By selecting a document from the "same" answer we can investigate further why this was the case.

3.3 Generating an Exhibit

Exhibit is implemented as an open source JavaScript library and there is no software to install; everything works on the browser. RAVE generates two files for creating an Exhibit: an HTML file that contains the layout of the elements in the web page and the data file in json format. For producing the json view, in a similar

```
{
  types: {
  'Item': {
  pluralLabel: 'Items'
  }
 }
,items: [
{ type: 'Item',
label: '4', worker: '4',
querytype: 'informational',
answer: 'B',
image: 'Selena_GomeztopB.png',
thumbnail: 'Selena_GomeztopB_tb.png',
},
{ type: 'Item',
label: '1', worker: '1',
querytype: 'informational',
answer: 'Same',
image: 'Selena_GomeztopB.png',
thumbnail: 'Selena_GomezopB_tb.png',
},
... ] }
```

Figure 3: Exhibit collection source code. A json snippet that describes the data set.

fashion to Pivot, the prototype reads the results of the experiment output and produces all the information needed.

Figure 3 describes sample data and Figure 4 shows the exploration of the results sets for a specific navigational query to investigate if there was any difference in the preferences for both rankers.

Figure 5: Analyzing workers' work units by quantity.

3.4 Work Data Exploration

Now that we have shown how to visualize a data set, we can explore workers' data to identify patterns. For example, of the nine workers who participated in the task, three of them have performed most of the work and two worked on single assignments (Figure 5). By zooming into the workload of a particular worker we can investigate his/her answers in more detail by comparing to the answer of other workers. The tool then allows the user to explore not only the final labels but also the specific worker data that can help determine worker quality and other behaviors.

4. RELATED WORK

As researchers and practitioners collect and analyze their own labeled data sets, new tools and solutions that facilitate such tasks are becoming available. Examples are end-to-end industrial crowdsourcing pipelines [3], the automation of crowdsourcing relevance with Terrier [6], and an open source system for collecting relevance assessments [4]. On the visual analytics front, VIRTUE, a system for exploring IR system performance and related metrics is described in [2]. SeeDB, a visualization recommendation engine for fast visual analysis is presented in [8]. Finally, there is emerging work on using visualization to help collect good labels via crowdsourcing in the NLP annotations [5].

5. CONCLUSION AND FUTURE WORK

We showed a prototype that can automatically generate a facet-based visualization for exploring a collection of relevance assessments collected via crowdsourcing. While this may look like a very narrow space, in practice, practitioners spend considerable amount of time looking at labeled data before the relevance modeling phase. Our goal is to assist data analysts who need to collect and assess relevance tasks labels by allowing them to visually explore those data sets in more detail. As an example, we showed an A-B comparison experiment but the techniques presented work for any type of task that requires workers to visually explore content and produce some label.

We are not interested in imposing a particular visualization metaphor but rather to suggest the adoption of this type of tools as part of the relevance assessment gathering process in IR. The prototype offers the visualization for two tools and can be extended to others. The Exhibit example is very flexible and easy to deploy making it a low-cost development alternative.

Visually exploring a data set can be useful to decide if the labels are of good quality, if there are no potential issues with the experiment or if the presentation of the results can bias the final labels. RAVE differs from previous research work in the sense that our focus is on exploring data sets instead of visualizing metrics. With RAVE the user can identify patterns and perform comparisons.

Future work includes automating the recommendation of visualizations, using the prototype to explore other existing assessments data sets like the TREC relevance labels and investigate the integration with other toolkits like D3.

6. REFERENCES

[1] Omar Alonso. Visualization for relevance assessments. *SIGIR Forum*, 48(2):14–21, 2014.

[2] Marco Angelini, Nicola Ferro, Giuseppe Santucci, and Gianmaria Silvello. VIRTUE: A visual tool for information retrieval performance evaluation and failure analysis. *J. Vis. Lang. Comput.*, 25(4):394–413, 2014.

[3] Vasilis Kandylas, Omar Alonso, Shiroy Choksey, Kedar Rudre, and Prashant Jaiswal. Automating crowdsourcing tasks in an industrial environment. In *HCOMP*, 2013.

[4] Bevan Koopman and Guido Zuccon. Relevation!: an open source system for information retrieval relevance assessment. In *SIGIR*, pages 1243–1244, 2014.

[5] Hanchuan Li, Haichen Shen, Shengliang Xu, and Congle Zhang. Visualizing NLP annotations for crowdsourcing. *CoRR*, abs/1508.06044, 2015.

[6] Richard McCreadie, Craig Macdonald, and Iadh Ounis. Crowdterrier: automatic crowdsourced relevance assessments with terrier. In *SIGIR*, page 1005, 2012.

[7] Tamara Munzner. *Visualization Analysis and Design*. A. K. Peters, 2014.

[8] Manasi Vartak, Samuel Madden, Aditya G. Parameswaran, and Neoklis Polyzotis. SEEDB: automatically generating query visualizations. *PVLDB*, 7(13):1581–1584, 2014.

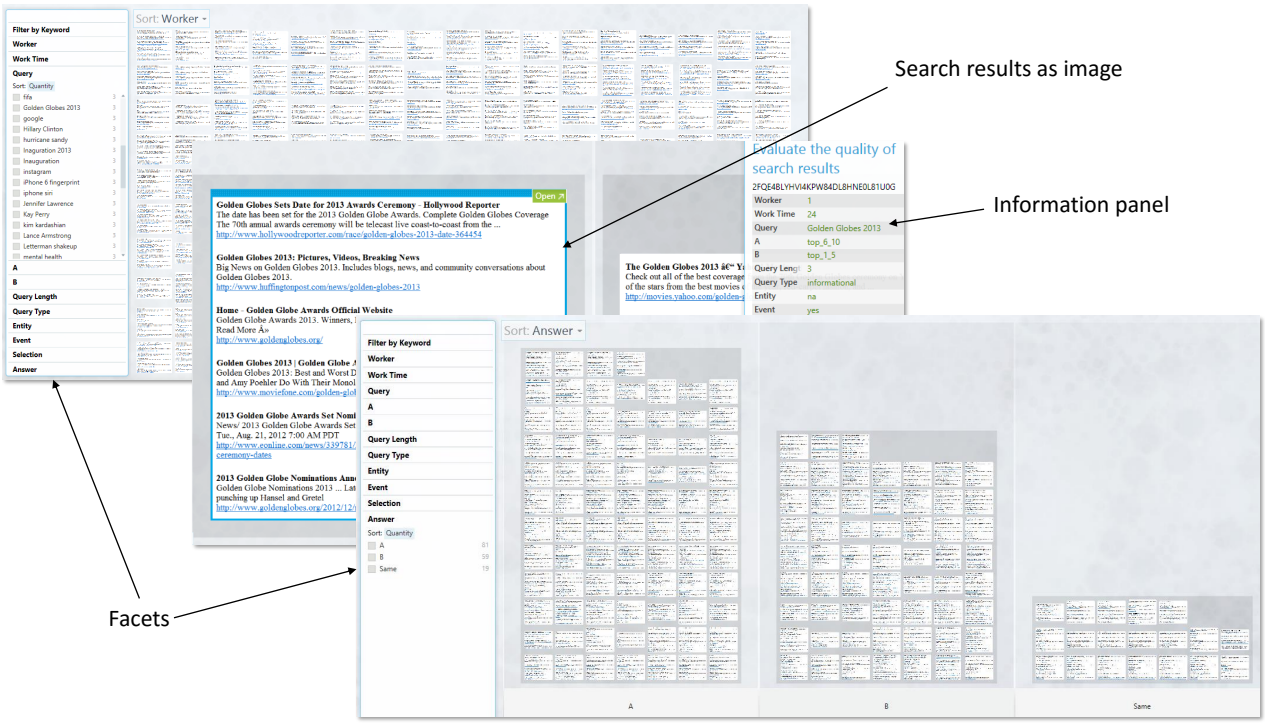

Figure 2: Pivot collection visualization. Three screenshots of the tool in action. From background to front: overview of the collection (all images) sorted by queries, focus on the search results for the query {Golden Globes 2013}, distribution of ranker preferences for a data set.

Figure 4: Exhibit collection visualization. Three screenshots of the tool in action. From background to front: overview of the collection with thumbnail images, focus on search results for the query {Selena Gomez}, more facets on the right side of the web page.